

A Bayesian Model of NMR Spectra for the Deconvolution and Quantification of Metabolites in Complex Biological Mixtures


*William Astle,  Maria De Iorio, [†]Sylvia Richardson,

**David Stephens and [‡]Timothy Ebbels.

*Department of Epidemiology, Biostatistics and Occupational Health,

**Department of Mathematics and Statistics,

McGill University, Montreal, Canada.

 Department of Statistical Science, University College, London,
United Kingdom.

[†]Department of Epidemiology and Biostatistics,

[‡]Section of Biomolecular Medicine, Department of Surgery and Cancer,
Faculty of Medicine, Imperial College, London. United Kingdom.

January 20, 2013

Authors' footnote

William Astle (E-mail: william.astle@mcgill.ca) is a Postdoctoral Fellow in the Department of Epidemiology, Biostatistics and Occupational Health, McGill University, Purvis Hall, 1020 Ave. des Pins Ouest, Montreal, QC, H3A 1A2, Canada. Maria De Iorio (E-mail: m.deiorio@ucl.ac.uk) is a Reader in Statistical Science in the Department of Statistical Science, University College, Gower Street, London, WC1E 6BT, UK. Sylvia Richardson (E-mail: sylvia.richardson@ic.ac.uk) is Professor of Biostatistics in the Department of Epidemiology and Biostatistics, Imperial College, St. Mary's Hospital Campus, London, W2 1PG, UK. David Stephens (E-mail: d.stephens@math.mcgill.ca) is Professor of Statistics in the Department of Mathematics and Statistics, McGill University, Montreal, QC, Canada, H3A 2K6. Timothy Ebbels (E-mail: t.ebbels@ic.ac.uk) is a Senior Lecturer in Computational Bioinformatics in the Section of Biomolecular Medicine, Department of Surgery and Cancer, South Kensington Campus, Imperial College, London. SW7 2AZ. This work was supported by UK BBSRC grant BB/E020372/1. Dr. Stephens is supported by a Discovery Grant from the Natural Sciences and Engineering Council of Canada. Dr. Astle's present position is supported by a Team Grant from the Fonds de recherche du Québec - Nature et technologies. The authors thank Jake Bundy for providing the yeast dataset and for help interpreting deconvolutions of the NMR spectra, Jie Hao for programming a C++ implementation and David Balding and Ernest Turro for comments on a draft manuscript.

Abstract

Nuclear Magnetic Resonance (NMR) spectra are widely used in metabolomics to obtain profiles of metabolites dissolved in biofluids such as cell supernatants. Methods for estimating metabolite concentrations from these spectra are presently confined to manual peak fitting and to binning procedures for integrating resonance peaks. Extensive information on the patterns of spectral resonance generated by human metabolites is now available in online databases. By incorporating this information into a Bayesian model we can deconvolve resonance peaks from a spectrum and obtain explicit concentration estimates for the corresponding metabolites. Spectral resonances that cannot be deconvolved in this way may also be of scientific interest so we model them jointly using wavelets.

We describe a Markov chain Monte Carlo algorithm which allows us to sample from the joint posterior distribution of the model parameters, using specifically designed block updates to improve mixing. The strong prior on resonance patterns allows the algorithm to identify peaks corresponding to particular metabolites automatically, eliminating the need for manual peak assignment.

We assess our method for peak alignment and concentration estimation. Except in cases when the target resonance signal is very weak, alignment is unbiased and precise. We compare the Bayesian concentration estimates to those obtained from a conventional numerical integration method and find that our point estimates have sixfold lower mean squared error.

Finally, we apply our method to a spectral dataset taken from an investigation of the metabolic response of yeast to recombinant protein expression. We estimate the concentrations of 26 metabolites and compare to manual quantification by five expert spectroscopists. We discuss the reason for discrepancies and the robustness of our methods concentration estimates.

KEYWORDS: metabolomics, concentration estimation, prior information, multi com-

⁴⁹ ponent model, block updates.

1. INTRODUCTION

Metabolomics (also known as *metabonomics* or sometimes *metabolic profiling*) is a scientific discipline concerned with the quantitative study of metabolites, the small molecules that participate in metabolic reactions. Research in this field is expanding rapidly, with applications in many areas of biology and medicine including cancer (e.g. Griffiths et al. (2002)), toxicology (e.g. Lindon et al. (2003)), organism classification (e.g. Bundy et al. (2002)), genetics (e.g. Illig et al. (2010)), biochemistry (e.g. Raamsdonk et al. (2001)), epidemiology (e.g. Holmes et al. (2008)) and disease diagnostics (e.g. Brindle et al. (2002)).

Almost all experiments in metabolomics rely on measurements of the abundances of metabolites in complex biological mixtures, often biofluid or tissue samples. One of the most extensively used techniques for obtaining such quantitative information is proton nuclear magnetic resonance (^1H NMR) spectroscopy. Metabolites generate characteristic resonance signatures in ^1H NMR spectra and each signature appears with intensity proportional to the concentration of the corresponding metabolite in the biological mixture.

Specialized models and tools are needed to draw inferences from ^1H NMR spectroscopic datasets, which are large and heavily structured. At present there is no statistical method for analyzing metabolomic NMR spectra reflecting the data generating mechanisms and the extensive prior knowledge available, e.g. on the form of metabolite NMR signatures. In this paper, we describe novel Bayesian approaches for the analysis of ^1H NMR data from complex biological mixtures. Specifically, we develop new models reflecting the data generation mechanisms and our prior knowledge, using a combination of parametric functions and wavelets. We introduce a computational strategy based on Markov chain Monte Carlo (MCMC), including novel block updates to overcome strong posterior correlation between the parametric functions

76 and the wavelets. We demonstrate the utility of the approach with simulations and
77 analyses of data from a yeast metabolomics experiment.

78 1.1 NMR spectroscopy

79 An NMR spectrum consists of a series of measurements of resonance intensity usu-
80 ally taken on a grid of equally spaced frequencies. Figure 1 is a representative sec-
81 tion (of about one tenth) of an NMR spectrum taken from an experiment into the
82 metabolomic response of yeast to recombinant protein expression. The x -axis of the
83 spectrum corresponds to resonant frequency and the y -axis to resonance intensity.

84 [Figure 1 about here.]

85 The spectrum is a collection of convolved peaks with different horizontal positions
86 and vertical scalings, each of which has the form of a Lorentzian curve. The zero
87 centered, standardized Lorentzian function takes the form

$$l_{\gamma}(x) = \frac{2}{\pi} \frac{\gamma}{4x^2 + \gamma^2}. \quad (1)$$

88 This is the pdf of a Cauchy distribution with scale parameter $\gamma/2$; in spectroscopy γ
89 is called the *peak-width at half-height* (or sometimes the *linewidth*).

90 Each spectral peak corresponds to magnetic nuclei resonating at a particular fre-
91 quency in the biological mixture. This frequency determines the displacement of the
92 peak on the x -axis, which is known as its *chemical shift* and is measured in parts
93 per million (ppm) of the resonant frequency of a standard peak. It is conventional
94 in NMR spectroscopy to use δ to denote chemical shift and for the δ -axis to increase
95 from right to left. ^1H NMR only detects the resonance of hydrogen nuclei and a
96 typical ^1H NMR spectrum has a range of about 0ppm-10ppm.

97 The resonant frequencies of a magnetic nucleus are largely determined by its
98 molecular environment, that is the chemical structure of the molecule in which it is

99 embedded and the configuration of its chemical bonding within the molecule. Con-
100 sequently, every metabolite has a characteristic molecular ^1H NMR *signature*, a con-
101 volution of Lorentzian peaks that appear in specific positions in ^1H NMR spectra.
102 These are the peaks observed in the ^1H NMR spectrum of a pure solution of the
103 metabolite. The peaks of a signature can have quite different chemical shifts (when
104 they are generated by protons with different bonding configurations) and so appear
105 widely separated in a spectrum.

106 Depending on its molecular environment, a proton may have more than one res-
107 onant frequency and when this happens the frequencies are usually very similar.
108 Consequently, the peaks generated appear in a spectrum as a juxtaposition called a
109 *multiplet*. The shape of a multiplet (number of peaks, their separations and relative
110 heights) can be used to identify the corresponding metabolite. Figure 2 shows an
111 ^1H NMR spectrum (top panel) and the resonance signatures of the four metabolites
112 contributing the principal resonance signals (lower panels), with characteristic peak
113 locations and multiplet shapes.

114 [Figure 2 about here.]

115 The intensity of a nuclear magnetic signal is proportional to the number of mag-
116 netically equivalent nuclei generating the resonance in the biological mixture. Conse-
117 quently, every resonance peak (and therefore also every metabolite resonance signa-
118 ture) scales vertically in a spectrum in proportion to the molecular abundance of the
119 corresponding compound in the mixture.

120 1.2 Specific challenges of NMR in metabolomics

121 Biofluids and tissue samples usually contain thousands of metabolites. However,
122 NMR is relatively insensitive, so that ordinarily a spectrum contains quantitative in-
123 formation on just a few hundred of the most abundant compounds. These compounds

124 can generate hundreds of resonance peaks in a spectrum, many of which overlap.

125 To quantify a collection of metabolites using NMR, at least one resonance peak
126 generated by each compound must be identified in the spectrum and deconvolved.
127 (To reduce uncertainty, it is desirable to identify as many peaks as possible for each
128 compound.) Estimates of the relative concentrations of the metabolites in the bio-
129 logical sample can be made by comparing the areas under the deconvolved resonance
130 peaks. (Estimates of absolute concentration require a reference compound).

131 The peak identification step (*assignment*) is complicated by fluctuations in peak
132 positions between spectra, caused by uncontrollable differences in experimental con-
133 ditions and differences in the chemical properties of the biological samples, such as
134 the pH and ionic strength. When this *positional noise* combines with peak overlap,
135 assignment can become very hard indeed.

136 [Figure 3 about here.]

137 The left panel of Figure 3 illustrates the problem. Excerpts from two spectra corre-
138 sponding to biological replicates from the same experiment are overlaid, focusing on a
139 doublet type multiplet with two peaks. The difference in peak position between repli-
140 cates is obvious to the eye. Here, the magnitude of the positional noise is insufficient to
141 confuse assignment by an expert spectroscopist but it will pose problems for standard
142 automated approaches. However, expert deconvolution is rarely practical because it
143 is labor intensive and relies on someone familiar with metabolite resonance patterns.
144 Targeted profiling (Weljie et al. 2006) against a standard library of metabolite reso-
145 nance peaks reduces the importance of expert knowledge but is slow because there is
146 no automated fitting procedure. Spectral binning (Holmes et al. (1994), Spraul et al.
147 (1994)) approaches divide the spectrum into regions (bins), within which the inten-
148 sity measurements are averaged, in an attempt to isolate distinct resonance signals.

149 Although this mitigates the effect of peaks fluctuating position within bins, fluctu-
150 ations across bin boundaries will cause anti-correlated increase/decrease of average
151 intensity in adjacent bins, even if there is no associated change in metabolite concen-
152 tration. Spectral binning balances parsimony and computational efficiency, retaining
153 quantitative information but in a representation with many fewer, easy to compute,
154 variables. However, the reduced variables are often analyzed without explicit quantifi-
155 cation of individual metabolites using pattern recognition methods such as principal
156 components analysis and partial least squares (Lindon, Holmes and Nicholson 2001).

157 The additional complication of positional noise combined with peak overlap is
158 illustrated in the right panel of Figure 3. The well defined resonance peaks overlap
159 with broad signals attributable to a combination of closely overlapping low metabolite
160 signals and/or macromolecular signatures. This introduces the problem of estimating
161 the proportion of the signal associated with the sharp resonances and the propor-
162 tion due to the broad component. The problem becomes even more complex when
163 the target peaks also overlap with other sharp metabolite signals which additionally
164 fluctuate between different spectra.

165 Currently, there is no statistical methodology that can simultaneously address the
166 problems of identification, deconvolution and quantification when there is positional
167 noise and peak overlap. We believe a method based on explicit quantification from
168 deconvolution of metabolite signatures should have significant advantages: spectral
169 convolution models are parsimonious because they correspond to the physical process
170 generating the data; the variables inferred are interpretable because they represent
171 concentrations of identified metabolites; these concentrations are of direct scientific
172 interest because they depend on the underlying biology.

1.3 Contributions of this paper

To tackle the problem of quantifying metabolites in complex biofluids such as cell supernatants or urine, we present a Bayesian model for ^1H NMR spectra and a Markov chain Monte Carlo (MCMC) algorithm to automate peak assignment and spectral deconvolution. Bayesian models for NMR data have been described before, notably in Bretthorst (1990*a*), in Bretthorst (1990*b*), in many subsequent papers by the same author, in Dou and Hodgson (1996) and in Rubtsov and Griffin (2007). Our modeling exploits extensive prior information on the resonance signatures of the metabolites, including the expected horizontal displacements and relative vertical scalings of the peaks. This novel approach allows us to deconvolve peaks and assign them to specific metabolites in a unified analysis, which eliminates the need for a manual assignment step. The prior information comes from the physical theory of NMR and from experimental information. Experimental resonance data on human metabolites are extensive and are publicly available, for example from the online database of the Human Metabolome Project, the HMDB (Wishart et al. 2009).

Almost all biofluid and tissue NMR spectra contain peaks for which there is no prior information in the presently incomplete public metabolite databases. Despite this, the component of a spectrum that cannot be assigned to known compounds, may contain metabolomic information that is scientifically useful, e.g. for classification of spectra. We therefore propose a two component joint model for a spectrum. We model the metabolites whose peaks we wish to assign explicitly parametrically, using information from the online databases, while we model the unassigned spectrum semi-parametrically, using wavelets. We choose wavelets because they model signal continuously but locally. They can account for the local correlation of a spectrum caused by the continuity of the underlying physical processes without imposing unrealistic global modeling constraints.

199 The wavelet component of our two component likelihood is extremely flexible
 200 so that, without restriction, it tends to absorb signal that should be modeled by
 201 the parametric component, thus inducing a lack of identifiability. We address this
 202 by penalizing the wavelet coefficients using heavy-tailed scale-mixture priors. These
 203 priors shrink wavelet coefficients wherever the spectral signal can be explained by
 204 the parametric component of the model. We also impose a truncation condition on
 205 the wavelets, which reflects prior knowledge that frequency-domain NMR spectra lie
 206 almost completely in the upper-half of the (x, y) plane.

207 To overcome the strong posterior correlation between parameters corresponding
 208 to the two model components we introduce purposely designed Metropolis-Hastings
 209 block proposals which update the parameters of the two components jointly.

210 2. MODELLING

211 2.1 NMR Spectra

212 Previous authors (Bretthorst (1990a), Dou and Hodgson (1996), Rubtsov and Griffin
 213 (2007)) developed Bayesian models for NMR data in the time domain, in which
 214 resonance signals appear as exponentially decaying sinusoids. However, we prefer
 215 to model conventionally preprocessed (by apodization, phase and baseline correction)
 216 data in the more interpretable frequency domain, in which resonance signals appear as
 217 peaks (e.g. Figure 1). Our model exploits the positivity of the frequency-spectrum, a
 218 condition which cannot be expressed parsimoniously in the time domain. Under an iid
 219 Gaussian model for errors, the two representations contain the same information since
 220 they are related by an orthogonal transformation (the discrete Fourier transform).

221 A frequency domain NMR dataset is a pair (\mathbf{x}, \mathbf{y}) , where \mathbf{x} is a length n vector
 222 of points on the chemical shift axis, usually regularly spaced and \mathbf{y} is a vector of
 223 corresponding resonance intensity measurements. n is typically of the order 10^3 to

224 10^4 depending on the resolution of the spectrum, and the size of the region under
 225 consideration. The intensity measurements are noisy, so that, although they measure
 226 inherently positive quantities, some components of \mathbf{y} are likely to fall below the δ -axis.
 227 \mathbf{y} is usually standardized in some way, for example so that $\sum_i^n y_i = 1$.

We model $\mathbf{y}|\mathbf{x}$ assuming the $y_i|\mathbf{x}$ are independent normal random variables and,

$$\text{E}(y_i|\mathbf{x}) = \phi(x_i) + \xi(x_i) \quad (2)$$

228 where the ϕ component of the model represents signal from metabolites with peaks
 229 we wish to assign explicitly and which have been previously characterized and cat-
 230 aloged in databases. (The metabolites chosen will vary from analysis to analysis
 231 according to the prior belief about the content of the biological mixture and the sci-
 232 entific question.) The ξ component of the model represents signal generated by peaks
 233 we do not wish to assign (this may include signal from uncataloged resonances of
 234 molecules which are partially characterized, with the characterized resonances mod-
 235 eled in the ϕ component). We construct ϕ parametrically (as a continuous function
 236 of continuous chemical shift δ) using the physical theory of NMR, and we model ξ
 237 semi-parametrically using wavelets.

238 2.2 Modeling ϕ , the cataloged metabolite signal

239 According to physics, the resonance signatures of distinct compounds are indepen-
 240 dent, accumulate with an intensity proportional to molecular abundance and aggre-
 241 gate in a spectrum by convolution. Consequently, we can write ϕ as a linear combina-
 242 tion, where each term in the sum corresponds to the signature of one of M different
 243 metabolites,

$$\phi(\delta) = \sum_{m=1}^M t_m(\delta)\beta_m. \quad (3)$$

(The value of M will depend on the scientific problem, it is likely to be of order 10^0 to 10^2). For each m , t_m is a continuous template function which specifies the NMR signature of metabolite m . The corresponding coefficient β_m is proportional to the molecular abundance of m (i.e. the concentration of m) in the biological sample. Physical theory restricts the model space for each template function to a parametric mixture of horizontally translated and vertically scaled Lorentzian peaks (Hore 1995).

We remarked previously that many metabolite signatures contain clusters of Lorentzian peaks called multiplets. Isolated peaks are often classed as *singlet* multiplets and, if we adopt this convention, we can express each signature template t_m completely as a linear combination of a set of multiplet curves g_{mu} ,

$$t_m(\delta) = \sum_u z_{mu} g_{mu}(\delta - \delta_{mu}^*), \quad (4)$$

where u is an index running over all the multiplets belonging to metabolite m . We assume $\int_0^\infty g_{mu}(\delta) d\delta = \int_{-\infty}^0 g_{mu}(\delta) d\delta$ for all m, u , so that the parameter δ_{mu}^* specifies the position on the chemical shift axis of the center of mass of the u th multiplet of the m th metabolite (see the large multiplet in Figure 4). We call δ_{mu}^* the chemical shift parameter of the multiplet. Each of the coefficients z_{mu} is a positive quantity, usually equal to the number of protons in a molecule of m that contribute resonance signal to the multiplet u . (z_{mu} may sometimes be non-integral, due for example to relaxation effects (Hore 1995). In these cases z_{mu} must be interpreted as an ‘effective’ proton contribution). $\int_{-\infty}^\infty g_{mu}(\delta) d\delta$ is constant over m and u so the area under each t_m is proportional to $\sum_u z_{mu}$, the number of protons resonating in a molecule of m .

With a few exceptions, multiplets can be classified into one of a number of common types (Figure 4) which determine the configuration of the peaks (a doublet, a triplet, a doublet of doublets, etc). This classification together with a small number of continuous quantities called J -coupling constants, which determine the (horizontal)

268 distances between the peaks, completely parameterize a multiplet curve.

269 [Figure 4 about here.]

To be precise, a multiplet curve g_{mu} is the weighted average of V_{mu} translated Lorentzian curves (see eqn. (1)),

$$g_{mu}(\delta) = \sum_{v=1}^{V_{mu}} w_{muv} l_{\gamma}(\delta - c_{muv}), \quad (5)$$

270 where the weights w_{muv} (which sum to one over v) determine the relative heights of
 271 the peaks of the multiplet and the translations c_{muv} determine the horizontal offsets
 272 of the peaks from the center of mass of the multiplet (see Figure 4). Multiplets are
 273 (usually) symmetric so that $\{-c_{muv} : v = 1, \dots, V_{mu}\} = \{c_{muv} : v = 1, \dots, V_{mu}\}$ and
 274 $w_{muv'} = w_{muv}$ when $c_{muv'} = -c_{muv}$.

275 2.3 Modeling ξ , the uncatalogued metabolite signal

276 We model $(\xi(x_1), \dots, \xi(x_n))^T$ as a linear combination of wavelet basis functions and
 277 use $\boldsymbol{\theta}$ to denote the vector of wavelet coefficients. We chose to use Daubechies's
 278 least asymmetric wavelets with 6 vanishing moments (symlet-6) as a wavelet basis
 279 because these wavelets have a similar shape to Lorentzian peaks. Symlets have been
 280 used previously to select features from NMR spectra (Kim et al. 2008) and sensitivity
 281 analysis comparing other potential wavelet bases showed little difference in spectral
 282 reconstructions.

283 2.4 The Likelihood

We now bring together the models for ϕ and ξ to make a formal specification of the probability model for the data. It is easier to do this in the wavelet domain, because the dimension of the wavelet space p often needs to be greater than the dimension of the data space n , to deal with distortion at the spectral borders (see Strang and

Nguyen (1996) and Section 1 of the supplementary material.) Let \mathcal{W} be the wavelet transform corresponding to the symlet-6 wavelet basis. The likelihood, is defined by

$$\mathcal{W}\mathbf{y} = \mathcal{W}\mathbf{T}\boldsymbol{\beta} + \mathbf{I}_p\boldsymbol{\theta} + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim N(0, \mathbf{I}_p/\lambda), \quad (6)$$

where \mathbf{T} is the $n \times M$ matrix with $t_m(x_i)$ as its (i, m) th entry, \mathbf{I}_p is the $p \times p$ identity matrix and where λ is a scalar precision parameter.

Equation (6) is a linear regression of $\mathcal{W}\mathbf{y}$ on the columns of $[\mathcal{W}\mathbf{T} \ \mathbf{I}_p]$, the matrix generated by adjoining $\mathcal{W}\mathbf{T}$ and \mathbf{I}_p columnwise. Since this matrix has more columns than rows, the regression coefficients cannot all be identifiable in the likelihood. We address this in the next section, by specifying a prior which helps to distinguish the parametric and semi-parametric components of the model.

2.5 Prior specification

Our aim is to obtain a joint Bayesian posterior distribution over the parameters controlling the shape of the templates $\{t_m : m = 1, \dots, M\}$, and the regression parameters $\boldsymbol{\beta}$, $\boldsymbol{\theta}$ and λ ; we now specify priors for these parameters.

Prior for the peak-width: Our focus is on spectra generated by biofluids such as cell supernatants or urine, for which peak-widths vary between, but negligibly within spectra; it is therefore reasonable to assume that peaks within a spectrum depend on a single common peak-width parameter γ . Our prior for γ is a log-normal distribution, with $\text{Median}(\gamma) = 1\text{Hz}/F$, $\text{Var}(\gamma) = 4.6\text{Hz}^2/F^2$ where F is the operating frequency of the spectrometer in MHz. This prior gives good support to a broad region around $1\text{Hz}/F$, typical of the peak-widths generated by modern spectrometers (Hore 1995). (With this prior, it is easy to relax the assumption of a common peak-width, since local deviations at the metabolite, multiplet, or peak level can be modelled using Gaussian random effects on $\log(\gamma)$.)

305 *Prior for the multiplets:* In section 2.2 we described a two-level parameterization of
 306 metabolite signature templates, defined by (4) and (5), as a linear combination of
 307 Lorentzian peaks nested in multiplets. This allows us to represent a difference in
 308 the uncertainty of peak positions within and between multiplets. The parameters
 309 $c_{mu\nu}$ and $w_{mu\nu}$, which determine the multiplet shapes, vary very little across NMR
 310 spectra. We assume they are constant and compute them by applying some simple
 311 rules (see Hore (1995), chap. 3 for the details), from empirical estimates of the J -
 312 coupling constants which are published in online databases. In contrast, as noted in
 313 section 1.2, the multiplet chemical shift parameters δ_{mu}^* do fluctuate slightly between
 314 spectra according to experimental conditions. We use an estimate $\hat{\delta}_{mu}^*$ of each δ_{mu}^* ,
 315 taken from online databases, to construct an informative prior which accounts for this
 316 uncertainty. The positional noise is local and smaller fluctuations are more probable
 317 than larger ones, so we assign each δ_{mu}^* a truncated normal prior distribution with
 318 mean parameter $\hat{\delta}_{mu}^*$, variance parameter 10^{-4}ppm^2 and truncation region $[\hat{\delta}_{mu}^* -$
 319 $0.03\text{ppm}, \hat{\delta}_{mu}^* + 0.03\text{ppm}]$. It may sometimes be appropriate to specify a multiplet
 320 or metabolite specific alternative depending on what is known about the variability
 321 of particular multiplet locations across spectra.

322 *Prior for the metabolite abundances:* Having defined a parametric prior for the
 323 metabolite signature templates, we now consider the prior for the vector β , each
 324 component of which represents the resonance intensity of a signature and is propor-
 325 tional to the abundance of that metabolite in the biological sample. The intensities
 326 are positive so the prior for each component of β should confine its support to \mathbb{R}^+ .
 327 For conditional conjugacy, we assign a normal prior to each β_m , truncated below at
 328 zero, $\beta_m \sim TN(e_m, 1/s_m, 0, \infty)$. This is flexible enough to encode prior information
 329 for a wide range of research problems. For the simulations and examples presented
 330 in this paper we assume low prior information and choose $e_m = 0$ and $s_m = 10^{-3}$ for

331 all m .

332 *Prior for the wavelet coefficients and precision parameter:* In section 2.4 we observed
 333 that the parametric (ϕ) and semi-parametric (ξ) components of the model are not
 334 identifiable in the likelihood. In order to resolve the model components we penalize the
 335 semi-parametric component by assigning the wavelet coefficients a prior distribution
 336 with heavy tails and a concentration of probability mass near zero. In addition, to
 337 reflect prior-knowledge that NMR spectra are mostly restricted to the half-plane above
 338 the chemical shift axis, our prior penalizes models in which $\mathcal{W}^{-1}\boldsymbol{\theta}$ has components
 339 below a small negative threshold. To be precise, in order to specify a joint prior
 340 for $(\boldsymbol{\theta}, \lambda)$ we introduce a vector of hyper-parameters $\boldsymbol{\psi}$, each component of which
 341 corresponds to a wavelet and a vector of hyper-parameters $\boldsymbol{\tau}$, each component of
 342 which corresponds to a spectral data point. The joint prior for $(\boldsymbol{\theta}, \boldsymbol{\psi}, \boldsymbol{\tau}, \lambda)$ has pdf
 343 proportional to

$$\lambda^{a+\frac{p+n}{2}-1} \left[\prod_{jk} \psi_{jk}^{c_j-1/2} \exp\left(-\frac{\psi_{jk} d_j}{2}\right) \right] \exp\left(-\frac{\lambda}{2} \left(b + \sum_{jk} \psi_{jk} \theta_{jk}^2 + r(\boldsymbol{\tau} - h\mathbf{1}_n)^2\right)\right) 1_{\{\mathcal{W}^{-1}\boldsymbol{\theta} \geq \boldsymbol{\tau} \wedge h\mathbf{1}_n \geq \boldsymbol{\tau}\}} \quad (7)$$

344 The index jk here corresponds to the k th wavelet in the j th wavelet-scaling level.

345 The following lemma is proved in the supplementary material

346 **Lemma 1** *Normalization of (7) defines a joint prior for $(\boldsymbol{\theta}, \boldsymbol{\psi}, \boldsymbol{\tau}, \lambda)$. This prior is*
 347 *proper.*

348 The joint prior specified by (7) was motivated by consideration of a scale mixture
 349 of multivariate normals with smoothed truncation limits:

$$P(\boldsymbol{\theta}|\boldsymbol{\psi}, \boldsymbol{\tau}, \lambda) = \frac{\lambda^{p/2} \prod_{jk} \psi_{jk}^{1/2}}{C_{\boldsymbol{\psi}\boldsymbol{\tau}\lambda}} \exp\left(-\frac{1}{2} \sum_{jk} \lambda \psi_{jk} \theta_{jk}^2\right) 1_{\{\mathcal{W}^{-1}\boldsymbol{\theta} \geq \boldsymbol{\tau}\}}, \quad (8)$$

$$\psi_{jk} \sim \text{Gamma}(c_j, d_j/2), \quad (9)$$

$$\tau_i \sim TN(h, 1/(\lambda r), -\infty, h), \quad (10)$$

$$\lambda \sim \text{Gamma}(a, b/2), \quad (11)$$

where $C_{\psi\tau\lambda}$ is a normalizing constant. The index i here corresponds to the i th spectral data point. In this specification, ψ allows the prior precision associated with each wavelet to deviate from the global precision λ . The gamma hyperprior on each component of ψ induces local shrinkage in the marginal prior for θ , which encourages posterior sparsity in the wavelet coefficients. τ is a vector of n truncation limits, which bounds $\mathcal{W}^{-1}\theta$ below. The decaying hyperpriors on the components of τ smooth these limits and penalize θ more heavily as more of the semi-parametric component (ξ component) of the model lies below the line $y = h$, where h is a small negative number, chosen close to zero on the spectral intensity scale.

The joint distribution specified by (8)-(10) is a reasonable representation of prior belief about θ and λ ; it places a constraint on the *conditional* distribution of θ given ψ , τ and λ . However, because the normalizing constant $C_{\lambda\psi\tau}$ of (8) has no closed form, it is hard to devise a computationally efficient scheme to sample from the resulting ‘doubly intractable’ posterior (Murray et al. 2006). The prior defined by (7) places the constraint on the *joint* distribution of the parameters, rather than the conditional distribution and it is easy to sample from the full conditionals of θ , ψ , τ and λ , if we use this prior. We contend that this is an equally valid specification and show in Section 2 of the supplementary material that it behaves similarly to the prior defined by (8)-(10).

Figure 5 demonstrates the effect of penalizing the semi-parametric component of the model when it lies below the chemical shift axis. First, without penalizing ξ in the lower half plane, if the components of θ are given untruncated Student’s- t priors, the posterior for the vector of quantification parameters β focuses asymptotically on a

373 region close to the ordinary least squares estimate of the parameters, $(\mathbf{T}^T \mathbf{T})^{-1} \mathbf{T}^T \mathbf{y}$,
 374 while the wavelet component absorbs most of the residual spectrum. When a metabo-
 375 lite has a multiplet embedded in a region of unassigned spectral resonance the least
 376 squares estimate overestimates the corresponding quantification parameter. The sig-
 377 nature templates absorb spectral signal even when they do not match the shape of the
 378 spectral data and this leads to strong posterior support for negative values for some
 379 components of $\mathcal{W}^{-1} \boldsymbol{\theta}$. The prior defined by (7) however, can correct the concentra-
 380 tion estimate providing at least one of the multiplets of the metabolite is deconvolved
 381 cleanly.

382 [Figure 5 about here.]

383 *Scale-invariant inference and hyper parameter settings:* Note that in the limit $a \rightarrow$
 384 $0, b \rightarrow 0, \forall m, s_m \rightarrow 0$, the (improper) prior is invariant under the scaling reparame-
 385 terisation $(\boldsymbol{\beta}, \boldsymbol{\theta}, \lambda) \mapsto (S\boldsymbol{\beta}, S\boldsymbol{\theta}, \lambda/S^2)$ for every constant $S > 0$ so that, in this limit,
 386 inference is unaffected by the scale of measurement of the data. Smaller a and b
 387 correspond to increased uncertainty in the value of λ . For simulations and examples
 388 described in this paper we take $a = 10^{-9}$ and $b = 10^{-6}$.

389 The values of the c_j and d_j control the degree of shrinkage penalization imposed on
 390 the wavelet coefficients. Experience shows $c_j = 0.05$ and $d_j = 10^{-8}$ provides adequate
 391 penalisation. More stringent penalization is possible but our MCMC algorithm tends
 392 to mix less well as the penalty gets stronger, because our block updates are less good
 393 at targeting the true posterior distribution of $\boldsymbol{\theta}$ (see below). h controls the threshold
 394 below which the wavelets are penalised in the lower half plane while r controls the
 395 strength of that penalisation. We choose h a little below zero to be sure that true
 396 signals below $y = 0$, due for example to baseline wiggle can be adequately modelled
 397 by wavelets. We set $h = -0.002$ and $r = 10^5$.

3. MARKOV CHAIN MONTE CARLO ALGORITHM

We implemented a Markov chain Monte Carlo algorithm, to sample from the joint posterior distribution of the model parameters. There are three types of MCMC update.

- Firstly, there are Gibbs samplers for the components of β (truncated normal), the components of θ (truncated normal), the components of ψ (gamma), the components of τ (truncated normal) and λ (gamma). The specific distributions for these updates are given in the supplementary material.
- Secondly, there is a Metropolis-Hastings update for each of the parameters (except the components of β) controlling the parametric component of the model. Specifically, in order to update the multiplet chemical shift parameter δ_{mu}^* we propose $\delta_{mu}^{* \prime}$ from

$$TN(\delta_{mu}^*, V_{\delta_{mu}^*}^2, \hat{\delta}_{mu}^* - 0.03\text{ppm}, \hat{\delta}_{mu}^* + 0.03\text{ppm}), \quad (12)$$

a Gaussian proposal, centred on the current parameter value and truncated at the boundaries of the prior distributions. We adapt the proposal variance $V_{\delta_{mu}^*}^2$ using the Adaptive Metropolis-within-Gibbs algorithm of Roberts and Rosenthal (2009). $V_{\delta_{mu}^*}$ is tuned to target an acceptance rate of 0.44 by increments and decrements which decay in magnitude asymptotically like the inverse of the square root of the iteration number.

In order to update the peak-width parameter we make a Gaussian proposal for $\log \gamma$, centred on the current parameter value. Again, we adapt the proposal variance following Roberts and Rosenthal (2009).

The likelihood constrains $\mathcal{W}\mathbf{y} - \mathcal{W}\mathbf{T}\beta - \theta$, inducing strong posterior correlation between the parametric and semi-parametric components of the model. Con-

sequently, updates of this type, which just propose changes to the parametric component, can only make local moves in the state space of the Markov chain.

- Thirdly, there are Metropolis-Hastings block updates in each of which a parameter controlling the parametric component of the model is updated jointly with the vector $\boldsymbol{\theta}$ of wavelet coefficients. The joint proposal breaks the posterior correlation between the parametric and semi-parametric model components, allowing larger jumps in state-space.

Block updates of the δ_{mu}^* extend the univariate proposals described previously. First we draw the univariate proposal $\delta_{mu}^{*\prime}$ (although we fix the proposal variance, see supplementary material), then conditional on the value drawn we propose a new value $\boldsymbol{\theta}'$ for the vector of wavelet coefficients $\boldsymbol{\theta}$ and perform a global Metropolis-Hastings accept-reject assessment for the block update.

The conditional proposal for $\boldsymbol{\theta}'$ is a multivariate truncated normal distribution with mean parameter $\mathcal{W}\mathbf{y} - \mathcal{W}\mathbf{T}'\boldsymbol{\beta}$ (where \mathbf{T}' is the template matrix updated to reflect the initial univariate proposal), precision parameter $\lambda\mathbf{I}_p$ and truncation $\mathcal{W}^{-1}\boldsymbol{\theta}' \geq \boldsymbol{\tau}$. We can simulate from this distribution by making the change of basis, $\boldsymbol{\eta}' = \mathcal{W}^{-1}\boldsymbol{\theta}'$; the components of $\boldsymbol{\eta}'$ are then independent univariate truncated normal distributions. This choice of conditional proposal is motivated by the full conditional of $\boldsymbol{\theta}$,

$$\boldsymbol{\theta}|\boldsymbol{\psi}, \boldsymbol{\tau}, \lambda, \mathbf{y} \sim TMVN(\mathcal{W}\mathbf{y} - \mathcal{W}\mathbf{T}'\boldsymbol{\beta}, (\mathbf{I}_p + \boldsymbol{\Psi})^{-1}/\lambda, \mathcal{W}^{-1}\boldsymbol{\theta} \geq \boldsymbol{\tau}) \quad (13)$$

which has a similar distribution but with a reduced precision. Unfortunately there is no easy way to sample from (13) because the truncation condition $\mathcal{W}^{-1}\boldsymbol{\theta} \geq \boldsymbol{\tau}$ induces a complex dependence structure between the components of $\boldsymbol{\theta}$.

442 The only proposals not yet described are those for the block updates of each
 443 component of β with θ . We propose β'_m from a Cauchy distribution, truncated
 444 below at zero. We center the Cauchy distribution on the point that maximizes
 445 the full conditional of β_m subject to $\theta = \mathbf{0}$ and to the truncation condition
 446 $\mathbf{y} - \mathbf{T}\beta > \tau$. This is a greedy proposal, in the sense that it attempts to maximise
 447 β'_m and explain as much of the spectral signal as possible using the template
 448 for metabolite m , excluding the wavelets ($\theta = \mathbf{0}$). Conditional on the proposed
 449 β'_m we propose θ' from a multivariate truncated normal distribution with mean
 450 parameter $\mathcal{W}\mathbf{y} - \mathcal{W}\mathbf{T}\beta'$ (where β' is β with the m th component set to β'_m),
 451 precision parameter $\lambda\mathbf{I}_p$ and truncation $\mathcal{W}^{-1}\theta' \geq \tau$.

452 Although, on average, block updates move the chain further than the correspond-
 453 ing single parameter updates, the acceptance rate is lower. Consequently moves of
 454 the second and third type make complementary contributions to MCMC mixing.

455 *Improving convergence and mixing:* During the burn in stage of the MCMC we tem-
 456 per the likelihood and penalize the wavelet component of the model stringently to help
 457 the chain move into a region of good posterior support. A parameter T (see section
 458 2 of supplementary material), which jointly controls the temperature/penalization,
 459 gradually cools according to a deterministic schedule, proportional to the complement
 460 of a Gaussian cdf.

461 To improve mixing, we implemented a population (multi-chain) version of our
 462 algorithm in which the MCMC operates on a product state-space composed of copies
 463 of the space described in Section 2. The chain targets a tempered version of the
 464 posterior distribution of interest marginally on each subspace, but with each subspace
 465 taking a different value of T . The Metropolis-Hastings updates already described
 466 operate within each subspace; the acceptance probabilities for these depend only on

467 the state of the relevant subchain and so they can be carried out in parallel on different
 468 CPU cores. Additionally, there are two types of update which allow the transfer of
 469 information between subspaces. Firstly, we propose that subspaces adjacent in the
 470 ordering of T swap parameter values, as in the exchange moves of parallel tempering
 471 (Geyer 1991). Secondly, for a multiplet location parameter δ_{mu}^* in a given subchain,
 472 we pick a complementary subchain uniformly at random and propose a new value
 473 for δ_{mu}^* from a Cauchy distribution, centred on the value of the parameter in the
 474 complementary chain, with scale-parameter 3ppm. We then propose a conditional
 475 update for the θ of the given subchain in the manner of the block-updates already
 476 described. This proposal is made for every chemical shift parameter of every subchain
 477 once for each iteration of the MCMC algorithm. It allows good values of a chemical
 478 shift parameter to spread through the population of chains. In principle additional
 479 information sharing moves (e.g. Evolutionary Monte Carlo crossover (Liang and
 480 Wong 2000)) could be added, but we find this ‘copy’ move to be sufficient. Ergodicity
 481 and the marginal convergence of the subchain with $T = 1$ to a stationary distribution
 482 equal to the target posterior is assured by theory (Jasra, Stephens and Holmes 2007).
 483 We combine annealing with parallel chains by running a ladder in which the ratios of
 484 the temperatures of the subchains are constant over time, and for which the target
 485 subchain (that with $T = 1$) cools according to the specified annealing schedule.

486 4. PERFORMANCE

487 We simulated 100 biofluid NMR spectra by convolving the standard library spec-
 488 tra for pure samples of acetic acid, alanine, betaine, creatine, glycine, glycolic acid,
 489 guanidoacetic acid, lactic acid, succinic acid, taurine, trimethylamine and trimethy-
 490 lamine oxide all of which generate resonances in the region 0ppm – 5ppm. The
 491 concentration of each metabolite was generated as a $U[0, 1]$ random variable and the
 492 multiplets of each metabolite were perturbed by draws from $U[-0.03\text{ppm}, 0.03\text{ppm}]$.

493 We mixed each simulated spectrum with a broad Gaussian pdf the mean of which
494 was drawn uniformly from the range of the spectrum (0ppm – 5ppm) and standard
495 deviation of which was fixed at 5.8ppm. This hump is typical of the broad baseline
496 distortions commonly generated by macromolecules in biofluid spectra.

497 Because each spectrum in the standard library was generated from a different ^1H
498 NMR experiment, the width of the peaks in a simulated spectrum can vary slightly
499 between resonances generated by different compounds. This variability can affect
500 inferences of concentration made with our model because of correlation between con-
501 centration and peak-width parameters. To deal with this artefact of the simulator,
502 we extended the model by adding random effects to allow for some (small) inter-
503 metabolite variability in the peak-width.

504 We applied the (single chain version of the) MCMC algorithm for the extended
505 model to each simulated spectrum, running for 5 000 iterations with a 3000 iteration
506 burn in.

507 4.1 Multiplet Localization

508 Figure 6 is a normalised histogram of the difference between the Bayesian poste-
509 rior mean estimate of each multiplet chemical shift and the true simulated chemical
510 shift, illustrating the quality of the peak localisation. The Bayesian posterior mean
511 estimates 85% of shifts within 0.002ppm and 90% within 0.015ppm.

512 [Figure 6 about here.]

513 4.2 Quantification

514 We compared our Bayesian method for quantification to the following algorithm for
515 numerical integration, which is slightly more sophisticated than conventional spectral
516 binning because it includes a peak identification step. To integrate a multiplet we used
517 the known simulated perturbation in chemical shift to identify the precise location of

the multiplet. In practice, this information would be unknown to a spectroscopist, at best he or she might be able to identify a multiplet by eye and choose bin boundaries based on the observed location. The information is, of course, unavailable to the Bayesian method. Using the pure compound spectrum we identified the region $[L, R]$ of the chemical shift axis that corresponds to the central 95% mass of the multiplet’s parametric template. Finally, we estimated the total area under the multiplet by $(\sum_{L < x_i < R} y_i) / (0.95N(L - R))$, where N is the number of x_i in $[L, R]$.

The plots in Figure 7 show the actual vs. estimated concentrations for estimates made by numerical integration and Bayesian posterior mean. The numerical method is systematically biased towards overestimation for two reasons. Firstly, it fails to take account of the Gaussian background hump and secondly it cannot distinguish resonance in the region $[L, R]$ generated by the metabolite of interest from other confounding signals. The Bayesian estimate is unaffected by the first problem because the background hump is modeled with wavelets; the second problem is dealt with by deconvolution. The mean quadratic error for the numerical estimator is 0.106 (equivalent to Pearson correlation 0.70), where the mean is taken over all simulation replicates and metabolites. In comparison, the quadratic error for the Bayesian posterior estimator is 0.017 (equivalent to Pearson correlation 0.89), over sixfold smaller.

[Figure 7 about here.]

To investigate the performance of the Bayesian method under different strengths of multiplet chemical shift perturbation we ran an additional 150 replicates of the simulations previously described in 3 blocks of 50 with perturbations drawn from $U[-0.01\text{ppm}, 0.01\text{ppm}]$, $U[-0.02\text{ppm}, 0.02\text{ppm}]$ and $U[-0.04\text{ppm}, 0.04\text{ppm}]$ in each block. We extended the truncation on the chemical shift priors to $\pm 0.045\text{ppm}$ the database estimate. The peak assignments and concentration estimation errors are

543 summarized in Table 1. There is no evidence of a major trend in the quality of peak
544 assignment or concentration estimation with the magnitude of multiplet perturbation.

545 [Table 1 about here.]

546 5. YEAST DATASET

547 To illustrate the performance of our method on a real dataset, we took three spectra
548 from the experiment investigating the metabolic response of the yeast *Pichia pastoris*
549 to recombinant protein expression (Tredwell, Edwards-Jones, Leak and Bundy 2011).
550 The spectra were generated using biological replicates prepared under the same con-
551 ditions. Consequently, the metabolic profiles of the samples are extremely similar
552 and the spectra contain essentially the same metabolite concentration information.
553 Nevertheless, the spectra are slightly different, because for example, of experiment
554 level positional noise in the chemical shifts of resonance peaks. By modeling the
555 three spectra jointly we can quantify metabolites using information from all three
556 replicates, while accounting for these experiment level differences.

557 We used the model described in Section 2 as a basis for a joint model of multiple
558 spectra. In the new model, the vector of metabolite quantification parameters β is
559 held in common across the spectra. All the remaining parameters are copied from
560 the original model, with a replicate set assigned to each spectrum. The MCMC
561 algorithm for the multiple-spectra model is very similar to the procedure described
562 for the original model. The Metropolis-Hastings updates involving components of β
563 need to be adjusted, to reflect the dependence on multiple spectra, but are similar
564 to those for the simpler model (see section 3 of the supplementary material). The
565 updates for the remaining parameters continue to be valid within each spectrum
566 because, conditional on β , the joint posterior factorizes into separate probability
567 models, each corresponding to a different spectrum.

568 Tredwell, Behrends, Geier, Liebeke and Bundy (2011) manually quantified 37
569 metabolites from these spectra, with each of the five authors assigning the resonances
570 and estimating concentrations independently. Not all the resonance patterns gener-
571 ated by the 37 metabolites take the form of the symmetric multiplets described in
572 Section 2. Multiplet shapes are sometimes distorted by strong interaction effects,
573 which we cannot easily include in our model because they are neither described by
574 a known parametric model nor cataloged in a public database. However, distorted
575 multiplets are still convolutions of Lorentzian peaks, so it is sometimes possible to
576 construct a template-based model by estimating the weights and translations of (5)
577 from an NMR spectrum of the relevant pure compound generated under similar ex-
578 perimental conditions. We were able to construct a parametric signature template
579 for 26 of the 37 compounds by combining public database information with param-
580 eter estimates from our laboratory library of pure compound spectra. However, we
581 were unable to construct complete signatures, containing a full complement of mul-
582 tiplets, for every metabolite. Although this precludes a complete deconvolution of
583 the spectral signal generated by compounds with unmodeled resonances, and the
584 omitted resonance signals will be absorbed into the wavelet component of the model.
585 Nevertheless, our main aim is to obtain accurate concentration estimates and this
586 is still achievable, providing at least one multiplet from each metabolite deconvolves
587 correctly.

588 We ran the MCMC procedure, with 8 parallel chains tempered on a ladder, for
589 20 000 iterations following a 10 000 iteration burn in. We made an adjustment
590 to the prior on the chemical shift parameters of the singlet multiplets generated
591 by Histidinol near to 7.25ppm and 8.19ppm by truncating the prior at ± 0.5 ppm
592 of the HMDB estimate rather than at ± 0.03 ppm. This represents prior knowledge
593 that the chemical shifts of those multiplets are more variable than is typical because

594 of sensitivity to chemical properties of the biofluid. Figure 8 shows the posterior
595 deconvolution of a heavily congested region (2.6ppm – 3.1ppm) and a region of broad
596 (0.8ppm – 1.4ppm) from one of the spectra.

597 The five spectroscopists deconvolved the spectra with the assistance of the widely
598 used Chenomx spectrographic software, which implements a form of targeted profil-
599 ing (Weljie et al. 2006). This is probably the most precise method currently available
600 for estimating metabolite concentrations from spectra, although its accuracy depends
601 on spectroscopists being able to make correct peak assignments. Figure 9 is a plot
602 showing the strong concordance between the concentration estimates of the spectro-
603 scopists and concentration estimates made by the Bayes posterior mean. Although
604 11 of the 26 posterior mean estimates lie outside the range of the spectroscopists’
605 estimates, there are only 3 cases of substantial discordance. The statistical estimates
606 for arginine and histidinol are substantially larger than the spectroscopist’s estimates
607 while that for malic acid is substantially lower. In the cases of arginine and histidinol
608 it is hard to fault the posterior deconvolution by visual inspection. The discrepancies
609 could be due to an error in the templates used by the spectroscopists to profile the
610 resonances or to experimental errors in the database estimates of the parameters used
611 to construct the metabolite signature templates of Bayesian model. In the case of
612 malate the discrepancy appears to have been caused by a multiplet misalignment. In
613 principle this could be resolved by adjusting the prior for that multiplet’s chemical
614 shift parameter in order to force the correct alignment.

615 [Figure 8 about here.]

616 [Figure 9 about here.]

6. DISCUSSION

Presently, automatic methods for analyzing biofluid NMR data rely on non-parametric pattern recognition techniques or are based on approximate numerical integration algorithms, such as spectral binning. These methods ignore a large amount prior information about the physical process generating the spectral data.

Prior information about a data generating process can easily be incorporated into a Bayesian analysis through specification of a likelihood and specification of a prior distribution for the parameters of the likelihood. We have shown that a Bayesian model for biofluid spectra, which exploits an informative parametric prior for the patterns of resonance generated by selected metabolites, can be used to deconvolve those resonances from a spectrum and to obtain explicit concentration estimates for the metabolites.

Simulations show that our MCMC algorithm usually identifies spectral resonance peaks precisely. Peak misalignment may occur when a target resonance for a multiplet of a template appears in the spectrum close to other stronger signals. The model may then encourage the template to align incorrectly with the stronger signals, even if they have the wrong shape. This is because the wavelet coefficients are heavily penalized in the prior but the parametric templates are not. Even when the model posterior concentrates around an incorrect deconvolution, the strong prior penalization on negative spectral signal means that posterior estimates of concentration can still be accurate, providing at least one multiplet for each metabolite deconvolves correctly. Concentration estimation, the main motivation for the modeling, is therefore quite robust to mis-assignment of spectral resonances.

It is worth noting that resonance mis-assignment is a problem for all methods, (including manual assignment by an expert; it is unavoidable for binning methods when peaks overlap) and our approach suggests two methods for resolving mistakes.

643 Firstly, signature templates corresponding to the compounds generating confounding
644 signals can be added to the parametric component of the model (providing they are
645 available). Secondly, the prior on the chemical shift parameter can be adjusted to fix
646 the position of a misaligned multiplet.

647 Our approach yields improved concentration estimates. A comparison with a
648 method for quantifying metabolites based on numerical integration shows the poste-
649 rior mean estimates of the Bayesian model to be 6 fold more accurate in quadratic
650 error, even when exact multiplet locations are given to the numerical integration
651 algorithm.

652 We are able to fit $M = 20$ metabolites to an $n = 3\,000$ datapoint spectrum in
653 about 60 minutes on a 2.2GHz desktop machine using less than 0.5GB of RAM when
654 running the MCMC algorithm for 10 000 iterations. (If multiple chains are required
655 to improve mixing, they will run in parallel on a multicore machine.) The number
656 of operations required by the MCMC algorithm is linear in n and M . The memory
657 requirement is linear in the number of MCMC iterations and quadratic in n . This rate
658 of computation can easily compete with the rate of acquisition of spectra by a typical
659 NMR laboratory, effectively removing a major bottleneck in laboratory pipelines.

660 An accurate, automatic method for estimating metabolite concentrations from
661 ^1H -NMR spectra will assist many research projects in metabolomics. The field relies
662 heavily on NMR for metabolite quantification and currently, even projects analyz-
663 ing a few tens of spectra use numerical integration for estimation. Bayesian model-
664 ing should become increasingly useful as prior information on metabolite resonance
665 patterns becomes more accurate and extensive. For example, with more detailed
666 information our template model could be extended to deal systematically with the
667 effects of interactions between multiplets. We plan to develop our model further and
668 to release an efficient implementation of our methodology capable of simultaneously

deconvolving the majority of metabolites assignable in the NMR spectra of complex biological mixtures.¹

REFERENCES

- Bretthorst, G. L. (1990a), “Bayesian Analysis. I. Parameter Estimation Using Quadrature NMR Models,” *Journal of Magnetic Resonance*, 88, 533–551.
- Bretthorst, G. L. (1990b), “Bayesian Analysis. II. Signal Detection and Model Selection,” *Journal of Magnetic Resonance*, 88, 552–570.
- Brindle, J., Antti, H., Holmes, E., Tranter, G., Nicholson, J., Bethell, H., Clarke, S., Schofield, P., McKilligin, E., Mosedale, D., and Grainger, D. J. (2002), “Rapid and noninvasive diagnosis of the presence and severity of coronary heart disease using ¹H-NMR-based metabonomics,” *Nature Medicine*, 8(12), 1439–1445.
- Bundy, J. G., Spurgeon, D. J., Svendsen, C., Hankard, P. K., Osborn, D., Lindon, J. C., and Nicholson, J. K. (2002), “Earthworm species of the genus *Eisenia* can be phenotypically differentiated by metabolic profiling,” *FEBS Letters*, 521(1-3), 115–20.
- Dou, L., and Hodgson, R. (1996), “Bayesian inference and Gibbs sampling in spectral analysis and parameter estimation: II,” *Inverse Problems*, 12, 121–137.
- Geyer, C. (1991), “Markov Chain Monte Carlo Maximum Likelihood,” in *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, ed. E. Keramigas, Fairfax VA: Fairfax: Interface Foundation, pp. 156–163.
- Griffiths, J. R., McSheehy, P. M. J., Robinson, S. P., Troy, H., Chung, Y.-L., Leek, R. D., Williams, K. J., Stratford, I. J., Harris, A. L., and Stubbs, M. (2002),

¹BATMAN, an R package based on a C++ implementation of our methodology, is available from URL: <http://www.ic.ac.uk/medicine/people/t.ebbels/>

691 “Metabolic changes detected by in vivo magnetic resonance studies of HEPA-1
692 wild-type tumors and tumors deficient in hypoxia-inducible factor-1beta (HIF-
693 1beta): evidence of an anabolic role for the HIF-1 pathway,” *Cancer Research*,
694 62(3), 688–95.

695 Holmes, E., Foxall, P., Nicholson, J., Neild, G., Brown, S., Beddell, C., Sweatman,
696 B., Rahr, E., Lindon, J., Spraul, M. et al. (1994), “Automatic data reduction
697 and pattern recognition methods for analysis of ^1H nuclear magnetic resonance
698 spectra of human urine from normal and pathological states,” *Analytical Bio-*
699 *chemistry*, 220(2), 284–296.

700 Holmes, E., Loo, R. L., Stamler, J., Bictash, M., Yap, I. K. S., Chan, Q., Ebbels,
701 T., De Iorio, M., Brown, I. J., Veselkov, K. A., Daviglus, M. L., Kesteloot,
702 H., Ueshima, H., Zhao, L., Nicholson, J. K., and Elliott, P. (2008), “Human
703 metabolic phenotype diversity and its association with diet and blood pressure.,”
704 *Nature*, 453(7193), 396–400.

705 Hore, P. (1995), *Nuclear Magnetic Resonance*, Vol. 32 of *Oxford Chemistry Primers*,
706 Oxford: Oxford University Press.

707 Illig, T., Gieger, C., Zhai, G., Römisch-Margl, W., Wang-Sattler, R., Prehn, C., Alt-
708 maier, E., Kastenmüller, G., Kato, B. S., Mewes, H.-W., Meitinger, T., de An-
709 gelis, M. H., Kronenberg, F., Soranzo, N., Wichmann, H.-E., Spector, T. D.,
710 Adamski, J., and Suhre, K. (2010), “A genome-wide perspective of genetic vari-
711 ation in human metabolism,” *Nature Genetics*, 42(2), 137–41.

712 Jasra, A., Stephens, D. A., and Holmes, C. C. (2007), “Population-based reversible
713 jump Markov chain Monte Carlo,” *Biometrika*, 94(4), 787–807.

- Kim, S. B., Wang, Z., Oraintara, S., Temiyasathit, C., and Wongsawat, Y. (2008),
 “Feature selection and classification of high-resolution NMR spectra in the complex wavelet transform domain,” *Chemometrics and Intelligent Laboratory Systems*, 90(2), 161–168.
- Liang, F., and Wong, W. (2000), “Evolutionary Monte Carlo: Applications to Cp model sampling and change point problem,” *Statistica Sinica*, .
- Lindon, J. C., Holmes, E., and Nicholson, J. K. (2001), “Pattern recognition methods and applications in biomedical magnetic resonance,” *Progress in Nuclear Magnetic Resonance Spectroscopy*, 39(1), 1–40.
- Lindon, J. C., Nicholson, J. K., Holmes, E., Antti, H., Bollard, M. E., Keun, H., Beckonert, O., Ebbels, T. M., Reily, M. D., Robertson, D., Stevens, G. J., Luke, P., Breau, A. P., Cantor, G. H., Bible, R. H., Niederhauser, U., Senn, H., Schlotterbeck, G., Sidelmann, U. G., Laursen, S. M., Tymiak, A., Car, B. D., Lehman-McKeeman, L., Colet, J. M., Loukaci, A., and Thomas, C. (2003),
 “Contemporary issues in toxicology the role of metabonomics in toxicology and its evaluation by the COMET project,” *Toxicology and Applied Pharmacology*, 187(3), 137–46.
- Murray, I., Ghahramani, Z., and MacKay, D. J. C. (2006), MCMC for Doubly-intractable Distributions,, in *UAI*, AUAI Press.
- Raamsdonk, L. M., Teusink, B., Broadhurst, D., Zhang, N., Hayes, A., Walsh, M. C., Berden, J. A., Brindle, K. M., Kell, D. B., Rowland, J. J., Westerhoff, H. V., van Dam, K., and Oliver, S. G. (2001), “A functional genomics strategy that uses metabolome data to reveal the phenotype of silent mutations,” *Nature Biotechnology*, 19(1), 45–50.

- 738 Roberts, G. O., and Rosenthal, J. S. (2009), “Examples of adaptive MCMC,” *Journal*
739 *of Computational and Graphical Statistics*, 18(2), 349–367.
- 740 Rubtsov, D. V., and Griffin, J. L. (2007), “Time-domain Bayesian detection and
741 estimation of noisy damped sinusoidal signals applied to NMR spectroscopy,”
742 *Journal of Magnetic Resonance*, 188(2), 367–79.
- 743 Spraul, M., Neidig, P., Klauck, U., Kessler, P., Holmes, E., Nicholson, J., Sweatman,
744 B., Salman, S., Farrant, R., Rahr, E. et al. (1994), “Automatic reduction of
745 NMR spectroscopic data for statistical and pattern recognition classification of
746 samples,” *Journal of Pharmaceutical and Biomedical Analysis*, 12(10), 1215–
747 1225.
- 748 Strang, G., and Nguyen, T. (1996), *Wavelets and Filter Banks*, Cambridge, MA:
749 Wellesley Cambridge Press.
- 750 Tredwell, G. D., Behrends, V., Geier, F. M., Liebeke, M., and Bundy, J. G. (2011),
751 “Between-Person Comparison of Metabolite Fitting for NMR-Based Quantita-
752 tive Metabolomics,” *Analytical Chemistry*, 83(22), 8683–8687.
- 753 Tredwell, G. D., Edwards-Jones, B., Leak, D. J., and Bundy, J. G. (2011), “The De-
754 velopment of Metabolomic Sampling Procedures for *Pichia pastoris*, and Baseline
755 Metabolome Data,” *PLoS One*, 6(1), e16286.
- 756 Weljie, A. M., Newton, J., Mercier, P., Carlson, E., and Slupsky, C. M. (2006),
757 “Targeted profiling: quantitative analysis of ^1H NMR metabolomics data,” *An-*
758 *alytical Chemistry*, 78(13), 4430–42.
- 759 Wishart, D. S., Knox, C., Guo, A. C., Eisner, R., Young, N., Gautam, B., Hau,
760 D. D., Psychogios, N., Dong, E., Bouatra, S., Mandal, R., Sinelnikov, I., Xia,
761 J., Jia, L., Cruz, J. A., Lim, E., Sobsey, C. A., Shrivastava, S., Huang, P., Liu,

762 P., Fang, L., Peng, J., Fradette, R., Cheng, D., Tzur, D., Clements, M., Lewis,
763 A., De Souza, A., Zuniga, A., Dawe, M., Xiong, Y., Clive, D., Greiner, R.,
764 Nazyrova, A., Shaykhutdinov, R., Li, L., Vogel, H. J., and Forsythe, I. (2009),
765 “HMDB: a knowledgebase for the human metabolome,” *Nucleic Acids Research*,
766 37(Database issue), D603–10.

767

List of Figures

768

1 Section from an NMR spectrum from a yeast experiment. 33

769

2 Decomposition of a metabolite NMR spectrum 34

770

3 Positional noise and peak overlap. 35

771

4 Some common multiplet types 36

772

5 The effect of a prior penalizing the ξ component in the lower half plane 37

773

6 Bayesian estimated multiplet shifts vs simulated shifts 38

774

7 Comparison of concentration estimates from simulated spectra. . . . 39

775

8 Deconvolutions of the yeast spectra 40

776

9 Comparison of concentration estimates 41




Figure 1: A section from an NMR spectrum from an experiment investigating protein expression in yeast. The x -axis measures chemical shift in parts per million (ppm). The y -axis measures relative resonance intensity.

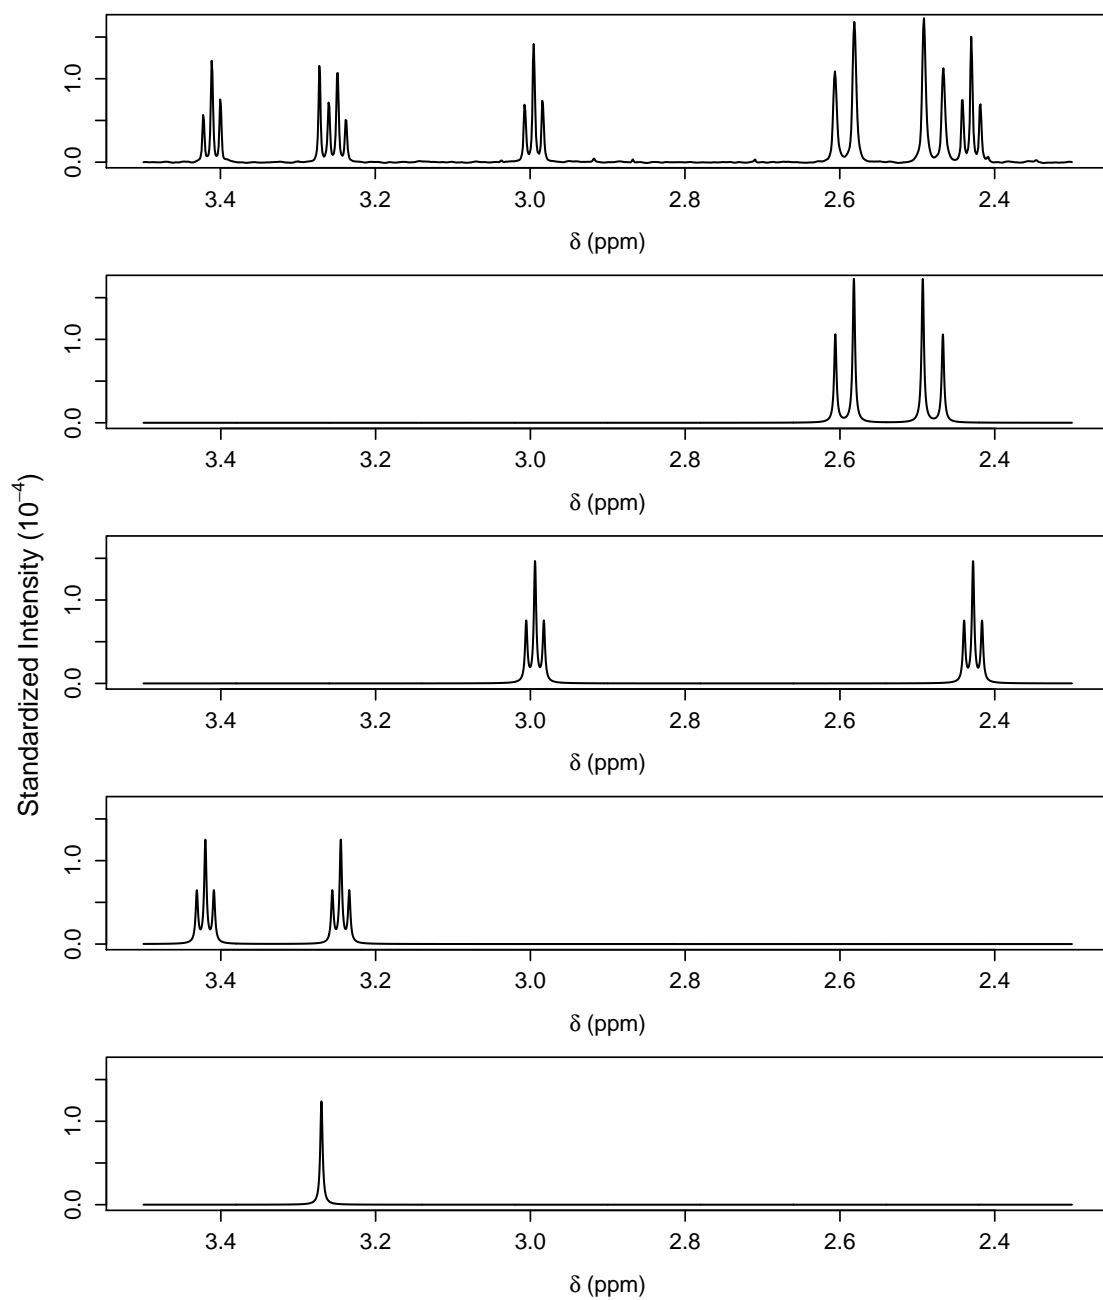


Figure 2: An ^1H NMR spectrum (*top panel*) with the principal resonance signals deconvolved into the metabolite NMR signatures (*lower panels*) of (in descending order by panel) citric acid, 2-oxoglutaric acid, taurine and trimethylamine N-oxide.




Figure 3: Positional noise between, and peak overlap within, two NMR spectra taken from the yeast experiment; resonances are generated by alanine (*left*) and threonine (*right*).




Figure 4: The peak configurations of some common types of multiplet. *Top row*: a doublet, with chemical shift δ_{mu}^* and peak offset $c_{mu\nu}$. *Middle row*: (from left to right) a singlet, a triplet, a quadruplet. *Bottom row*: (from left to right) a quintuplet, a doublet of doublets and a triplet of doublets.

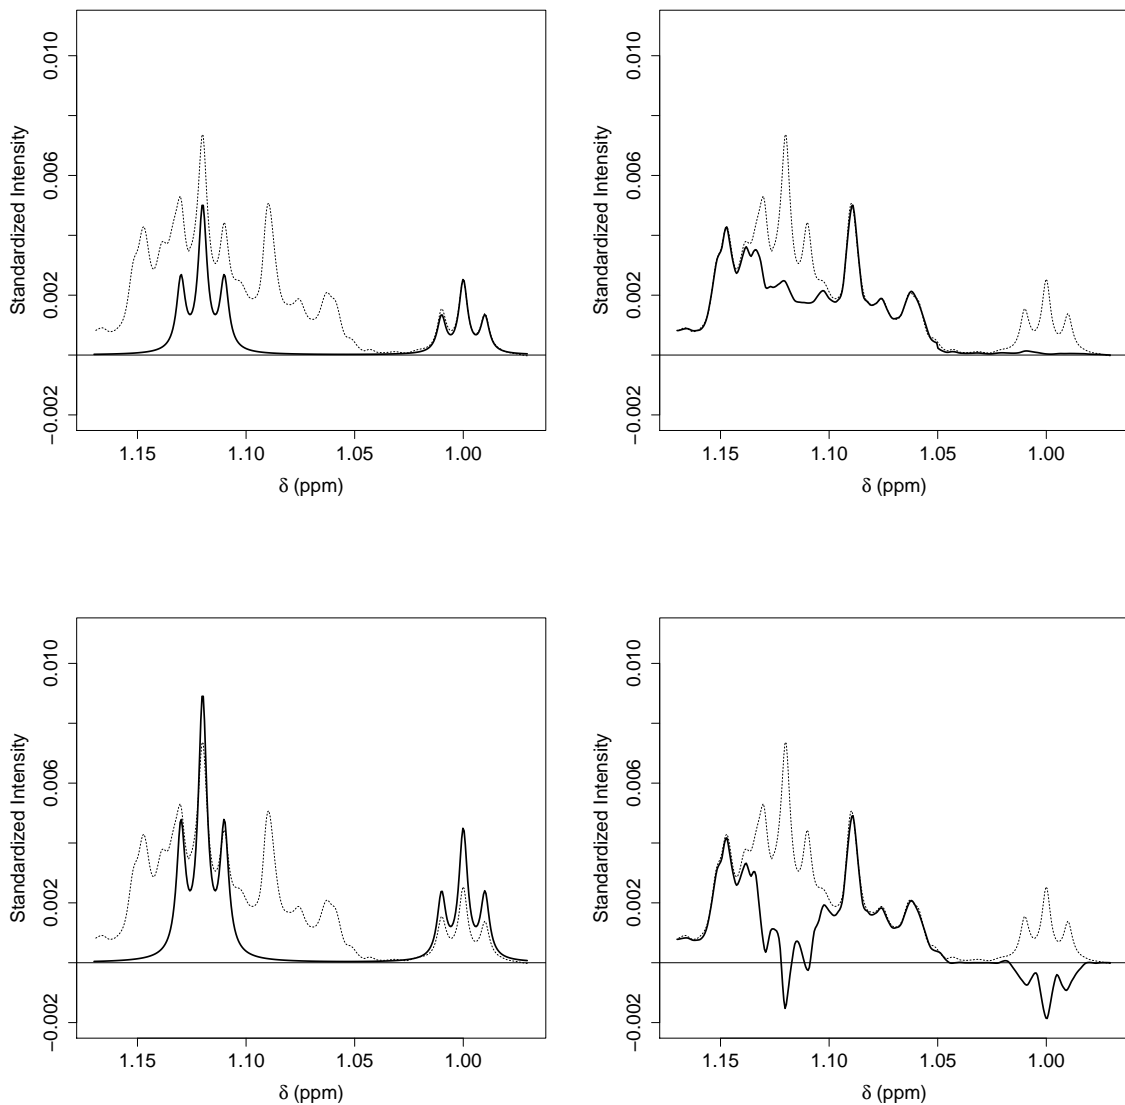


Figure 5: The effect of a prior penalizing the ξ component of the likelihood in the lower half plane (*top*) compared to one without this penalization (*bottom*). The dashed lines show the spectral data. Deconvolution of a parametric metabolite signature template (heavy lines, *left*) can be more accurate when the wavelet component (heavy lines, *right*) is penalized below the δ -axis.




Figure 6: Normalized histogram of simulated minus Bayesian posterior mean estimates of multiplet shift.

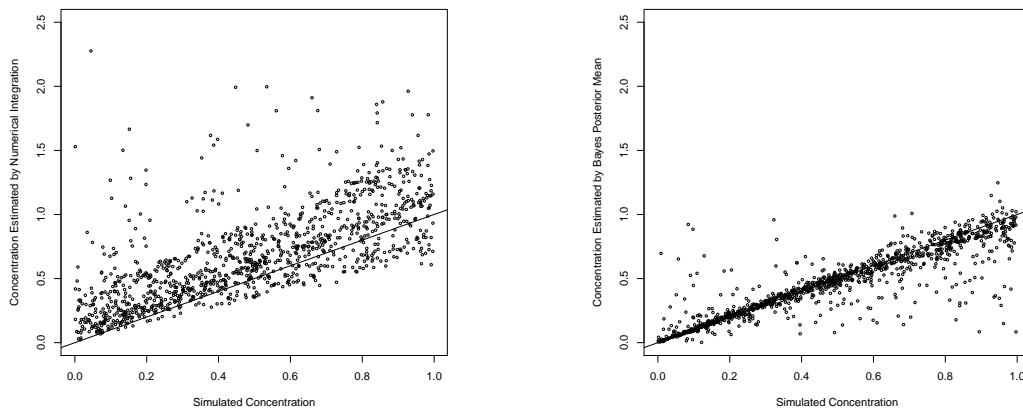



Figure 7: Concentration estimates from simulated spectra: true vs estimated concentrations for numerical integration (*left*) and for Bayesian posterior mean (*right*).






Figure 9: Comparison of metabolite concentration estimates by posterior mean of the quantification parameters from the Bayesian model (y -axis) with manual spectroscopists estimates (x -axis). The scales are calibrated so that the concentration of Trehalose is equal to 1. Each circle represents the mean of the five spectroscopists estimates while the limits of the horizontal bars represent the range. The inset is a magnification of the region bounded by the dashed line.

777
778
779

List of Tables

1	Peak assignment and concentration estimation errors for different magnitudes of simulated chemical shift perturbation.	43
---	--	----

Table 1: Peak assignment and concentration estimation errors for different magnitudes of simulated chemical shift perturbation.

Distribution of Simulated Perturbations	Peaks Assigned Within 0.002ppm	Mean Quadratic Error for Bayes Estimate	Mean Quadratic Error for Numerical Integration Estimate
$U[-0.01\text{ppm}, 0.01\text{ppm}]$	83%	0.016	0.141
$U[-0.02\text{ppm}, 0.02\text{ppm}]$	84%	0.016	0.102
$U[-0.04\text{ppm}, 0.04\text{ppm}]$	83%	0.017	0.135